QUESTION 41
When an OSPF design is planned, which implementation can help a router not have memory resource issues? Select the best response.

A. Have a backbone area (area 0) with 40 routers and use default routes to reach external destinations.
B. Have a backbone area (area 0) with 4 routers and 30,000 external routes injected into OSPF.
C. Have less OSPF areas to reduce the need for interarea route summarizations.
D. Have multiple OSPF processes on each OSPF router. Example, router ospf 1, router ospf 2

Answer: A
Explanation:
Memory issues usually come up when too many external routes are injected in the OSPF domain. A backbone area with 40 routers and a default route to the outside world would have less memory issues compared with a backbone area with 4 routers and 33,000 external routes being injected into OSPF. Router memory could also be conserved by using a good OSPF design. Summarization at the area border routers and use of stub areas could further minimize the number of routes exchanged.
The total memory used by OSPF is the sum of the memory used in the routing table (show ip route summary) and the memory used in the LSDB. The following numbers are a "rule of thumb" estimate. Each entry in the routing table will consume between approximately 200 and 280 bytes plus 44 bytes per extra path. Each LSA will consume a 100 byte overhead plus the size of the actual LSA, possibly another 60 to 100 bytes (For router links, this depends on the number of interfaces on the router). These amounts should be added to memory already used by other processes and by the IOS itself.
If you really want to know the exact number, you can do a show memory with and without OSPF being turned on. The difference in the processor memory used would be the answer.

QUESTION 42
LAB 1 - EIGRP Stub

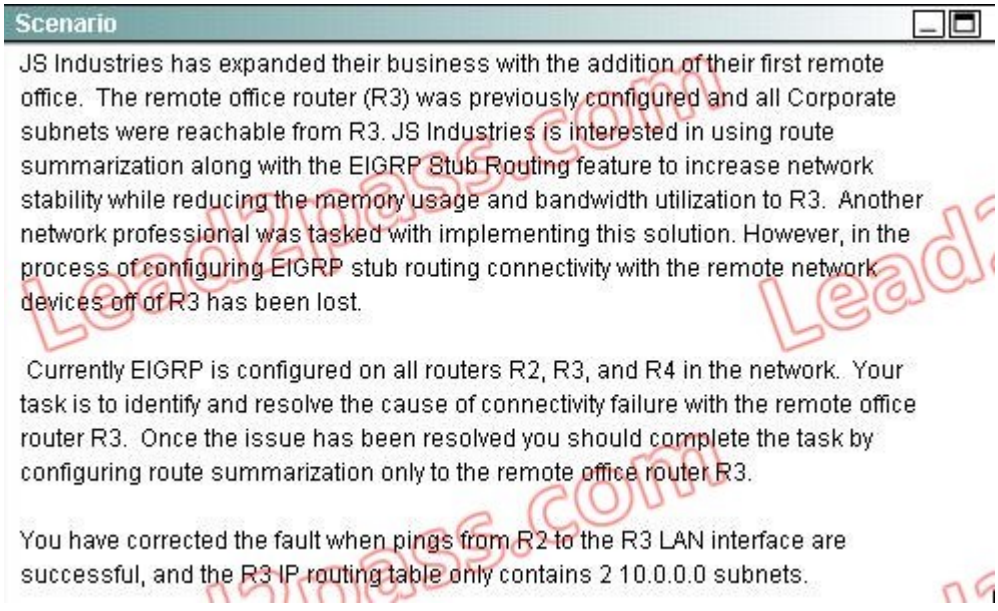

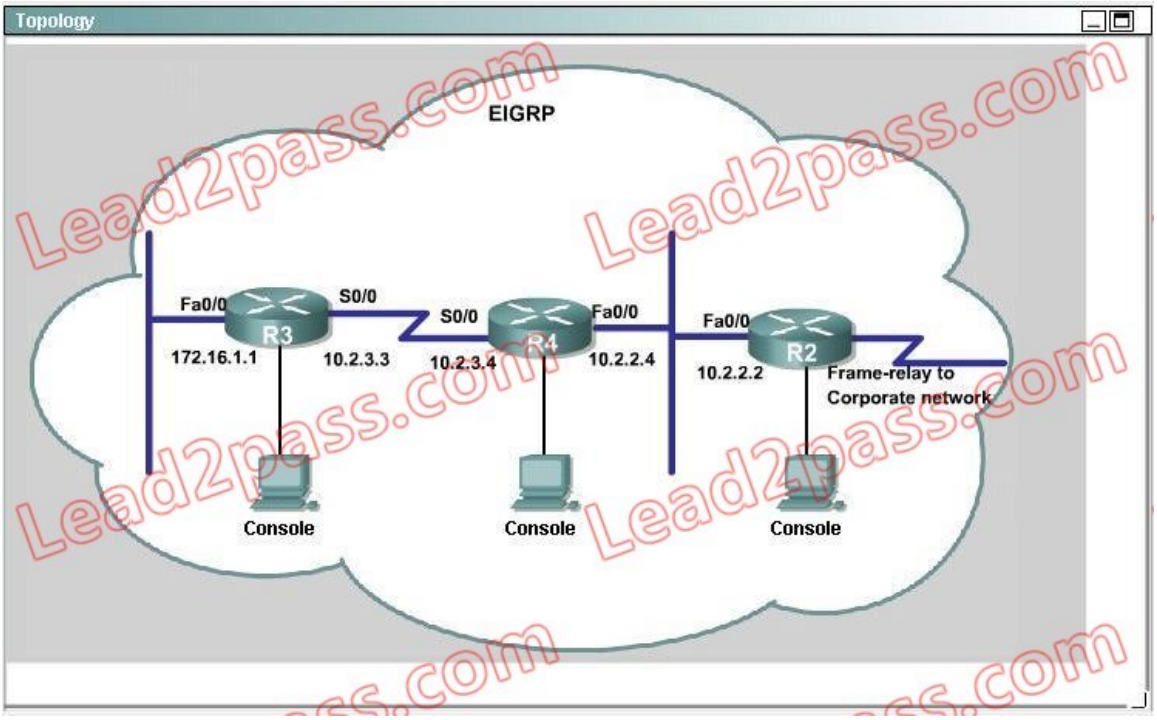

Answer and Explanation:
First we have to figure out why R3 and R4 can not communicate with each other. Use the **"show running-config"** command on router R3

```
R3#
R3#show run
Building configuration...

Current configuration : 731 bytes
!
 ip address 172.16.1.1 255.255.255.0
!
!
interface FastEthernet0/1
 no ip address
 shutdown
!
!
!
!
router eigrp 123
 network 10.0.0.0
 network 172.16.0.0
 no auto-summary
 eigrp stub receive-only
!
!
```

Notice that R3 is configured as a stub **receive-only** router. The receive-only keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system. This keyword will also prevent any type of route from being sent.

Therefore we will remove this command and replace it with the **eigrp stub** command:

```
R3# configure terminal
R3(config)# router eigrp 123
R3(config-router)# no eigrp stub receive-only
R3(config-router)# eigrp stub
R3(config-router)# end
```

Now R3 will send updates containing its connected and summary routes to other routers. Notice that the **eigrp stub** command equals to the **eigrp stub connected summary** because the **connected** and **summary** options are enabled by default.

Next we will configure router R3 so that it has only 2 subnets of 10.0.0.0 network. Use the **show ip route** command on R3 to view its routing table

```
R3# show ip route
```

```
Router3# show ip route

10.0.0.8/ is variably subnetted, 9 subnets, 2 masks

D    10.2.2.0/24 [90/30720] via 10.2.3.4, 00:00:06, Serial0/0
C    10.2.3.0/24 is directly connected, Serial0/1
D    10.2.4.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D    10.2.5.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D    10.2.6.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D    10.2.7.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D    10.2.8.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
D    10.2.9.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0

  172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks

D    172.16.0.0/16 is a summary, 02:04:27, Null0
C    172.16.1.0/24 is directly connected, FastEthernet0/0
```

Because we want the routing table of R3 only have 2 subnets so we have to summary sub-networks at the interface which is connected with R3, the s0/0 interface of R4.

There is one interesting thing about the output of the **show ip route** shown above: the **10.2.3.0/24**, which is a directly connected network of R3. We can't get rid of it in the routing table no matter what technique we use to summary the networks. Therefore, to make the routing table of R3 has only 2 subnets we have to summary other subnets into one subnet.

In the output if we don't see the summary line (like 10.0.0.0/8 is a summary…) then we should use the command ip summary-address eigrp 123 10.2.0.0 255.255.0.0 so that all the ping can work well.

In conclusion, we will use the **ip summary-address eigrp 123 10.2.0.0 255.255.0.0** at the interface s0/0 of R4 to summary.

```
R4> enable
R4# configure terminal
R4(config)# interface s0/0
R4(config-if)# ip summary-address eigrp 123 10.2.0.0 255.255.0.0
```

Now we jump back to R3 and use the **show ip route** command to verify the effect, the output is shown below:

```
Router3# show ip route

10.0.0.8/ is variably subnetted, 2 subnets, 2 masks

D    10.2.2.0/16 [90/2172416] via 10.2.3.4, 00:00:011, Serial0/0
C    10.2.3.0/24 is directly connected, Serial0/1

  172.16.0.0/24 is variably subnetted, 1 subnets

C    172.16.1.0/24 is directly connected, FastEthernet0/0
```


(But please notice that the ip addresses and the subnet masks in your real exam might be different so you might use different ones to solve this question)

But in your real exam, if you see the line "10.0.0.0/8 is a summary,….Null0" then you need to summary using the network 10.0.0.0/8 with the command "**ip summary-address eigrp 123 10.0.0.0 255.0.0.0**". This configuration is less optimize than the first but it summaries into 2 subnets as the question requires.

Finally don't forget to use the **copy running-config startup-config** command on routers R3 and R4 to save the configurations. (just skip if it's doesn't work.)

```
R4(config-if)# end
R4# copy running-config startup-config
```

QUESTION 43
Refer to the exhibit. A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP. The show command output shows that traffic will not flow between the networks. As a network consultant, you were asked to modify the configuration and certify the interoperability of the two networks. For traffic to flow from subnet 172.16.8.0/24 to the 172.16.16.0/24 subnet, which configuration change do you recommend?

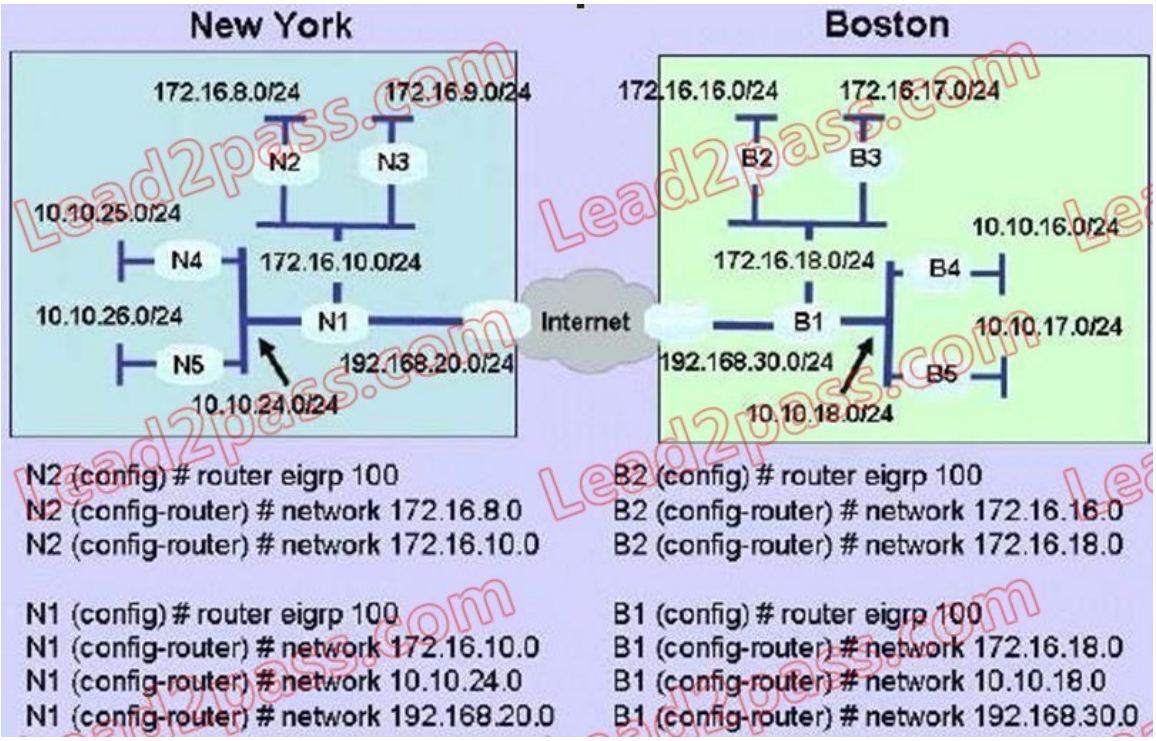

A. Turn off autosummarization on routers N1 and B1.
B. Add IP summary addresses to the Internet-pointing interfaces of routers N1 and B1.
C. Turn off autosummarization on routers N2 and B2.
D. Add wildcard masks to the network commands on routers N2 and B2.

Answer: A
Explanation:
Basically auto route summarization happens at the classful network boundary...so that would make N1 and B1 the locations that summarization would occur for the 172.16.0.0/16 classful networks.
So if you left auto-summarization enabled on those 2 routers, you would have an issue with discontiguous networks being advertised by both routers N1 and B1 with their classful mask (172.16.0.0/16 and 10.0.0.0/8), which will cause you issues.
Turning off auto-summarization on N2 and B2 wouldn't make any difference, as their networks wouldn't be summarized due to the fact that they are not meeting a classful boundary on their perspective routers. N1 will receive the 172.16.8.0/24 network from N2 with auto-summarization enabled.

QUESTION 44
Refer to the exhibit. A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP over EoMPLS. As a network consultant, you were asked to verify the interoperability of the two networks.
From the show ip route command output, what can you tell the customer about the traffic flow between the subnet in New York (172.16.8.0/24) and the subnets in Boston (172.16.16.0/24 and 10.10.16.0/24)? Select the best response.

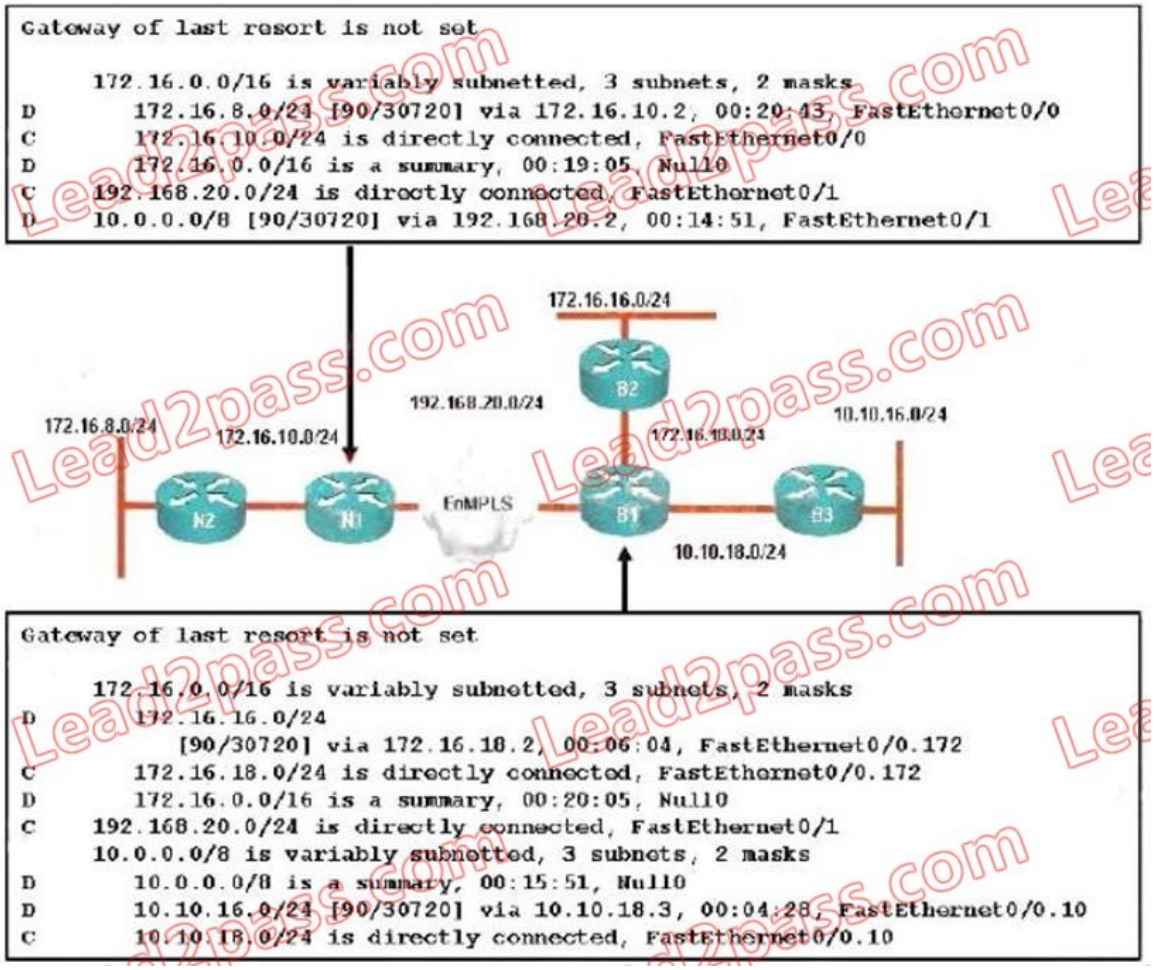

A. Traffic is flowing between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0 and no configuration changes are needed.
B. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0.
C. Traffic will flow between the 172.16.8.0 subnet and 172.16.16.0 without any further configuration changes. However, auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 10.10.16.0 subnet.
D. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 172.16.16.0 subnet. However, traffic will flow between the 172.16.8.0 subnet and 10.10.16.0 without any further configuration changes.

Answer: B
Explanation:
Basically auto route summarization happens at the classful network boundary...so that would make N1 and B1 the locations that summarization would occur for the 172.16.0.0/16 classful networks.
So if you left auto-summarization enabled on those 2 routers, you would have an issue with discontiguous networks being advertised by both routers N1 and B1 with their classful mask (172.16.0.0/16 and 10.0.0.0/8), which will cause you issues.
Turning off auto-summarization on N2 and B2 wouldn't make any difference, as their networks

wouldn't be summarized due to the fact that they are not meeting a classful boundary on their perspective routers.

QUESTION 45
Refer to the exhibit. You are the network administrator responsible for the NProuter, the 10.1.1.1 router, and the 10.1.1.2 router.
What can you determine about the OSPF operations from the debug output? Select the best response.

A. The NProuter has two OSPF neighbors in the "Full" adjacency state.
B. The NProuter serial0/0 interface has the OSPF dead timer set to 10 seconds.
C. The NProuter serial0/0 interface has been configured with an OSPF network type of "point- to-point".
D. The 10.1.1.1 and 10.1.1.2 routers are not using the default OSPF dead and hello timers setting.
E. The "Mismatched" error is caused by the expiration of the OSPF timers.

Answer: B
Explanation:
First we should understand clearly about the line
Dead R 120 C 10, Hello R 30 C 30
The "R" here means "Received" and "C" means "Configured". In other words, "Dead R" is the Dead Timer Received from the neighbor and the "Dead C" is the Dead Timer of the local router. Therefore in this case "Dead R 120 C 10 means the Death Timer of the neighbor is 120 seconds while the local Dead Timer is 10 seconds, which causes a mismatch. Also we can learn that the local OSPF dead timer is set to 10 seconds.
For your information, by default, OSPF uses a 10-second hello timer and 40-second hold timer on broadcast and point-to-point links, and a 30-second hello timer and 120-second hold timer for all other network types.

QUESTION 46
The maximum number of routers per OSPF area typically depends on which three factors? (Choose three.)

A. the kind of OSPF areas being implemented
B. the number of external LSAs in the network

C. the number of DRs and BDRs in the areas
D. the number of virtual links in the areas
E. how well the areas can be summarized
F. the use of LSA filters

Answer: ABE

QUESTION 47
Refer to the exhibit. You have completed an OSPF implementation, and you are verifying OSPF operation. You notice that router A and router B are stuck in the two-way state. From the show ip ospf interface command output, what is the cause of this issue?

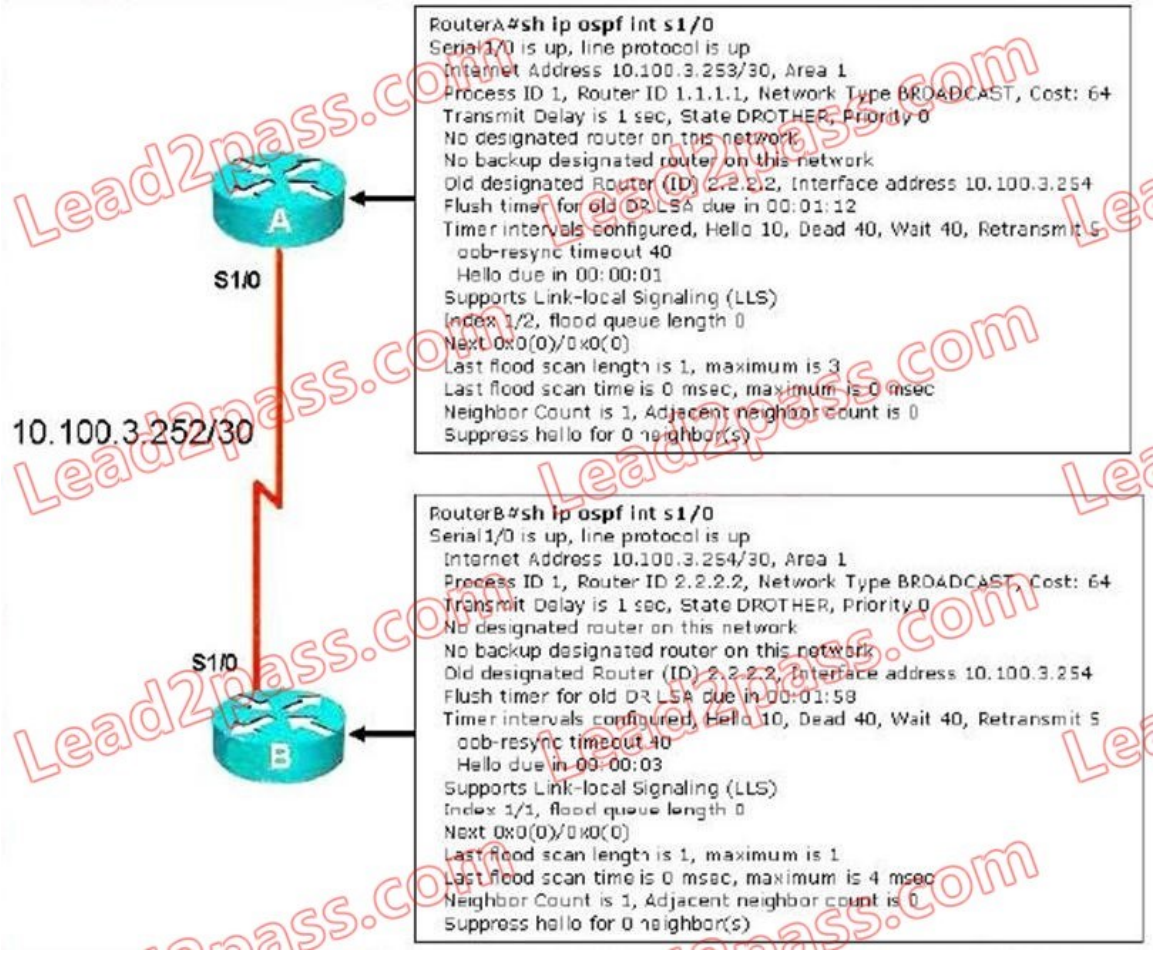

A. All OSPF implementations must have at least one interface in area 0.
B. You are attempting to run in the broadcast mode over an NBMA interface.
C. Both routers are configured to function as a BDR; therefore, there is no DR router.
D. Someone has changed the OSPF router ID; therefore you must clear the OSPF process.
E. The OSPF priority is set to 0 on both routers; therefore neither can become the DR.

Answer: E
Explanation:
When OSPF adjacency is formed, a router goes through several state changes before it becomes

fully adjacent with its neighbor. The states are Down, Attempt, Init, 2-Way, Exstart, Exchange, Loading, and Full.
An OSPF neighbor reaches the 2-way state when bidirectional communication is established (each router has seen the other's hello packet). This is the beginning of an OSPF adjacency. On broadcast media and non-broadcast multiaccess networks, the DR and BDR are elected in this state. But the priority on both routers are 0 so no DR and BDR are elected -> These routers stay in the 2-way state.

QUESTION 48
You are troubleshooting an OSPF problem where external routes are not showing up in the OSPF database Which two options are valid checks that should be performed first to verify proper OSPF operation? (Choose two.)

A. Are the ASBRs trying to redistribute the external routes into a totally stubby area?
B. Are the ABRs configured with stubby areas?
C. Is the subnets keyword being used with the redistribution command?
D. Is backbone area (area 0) contiguous?
E. Is the CPU utilization of the routers high?

Answer: AC
Explanation:
A totally stubby stubby area cannot have an ASBR so it will discard this type of LSA (LSA Type 5) -> A is a valid check.
Each stubby area needs an ABR to communicate with other areas so it is normal -> B is not a valid check.
When pulling routes into OSPF, we need to use the keyword "subnets" so that subnets will be redistributed too. For example, if we redistribute these EIGRP routes into OSPF:
+ 10.0.0.0/8+ 10.10.0.0/16+ 10.10.1.0/24
without the keyword "subnets"
router ospf 1redistribute eigrp 1
Then only 10.0.0.0/8 network will be redistributed because other routes are not classful routes, they are subnets. To redistribute subnets we must use the keyword "subnets" router ospf 1redistribute eigrp 1 subnets
-> C is a valid check.
We don't need to care if area 0 is contiguous or not -> D is not a valid check. CPU utilization cannot be the cause for this problem -> E is not a valid check.

QUESTION 49
When verifying the OSPF link state database, which type of LSAs should you expect to see within the different OSPF area types? (Choose three.)

A. All OSPF routers in stubby areas can have type 3 LSAs in their database.
B. All OSPF routers in stubby areas can have type 7 LSAs in their database.
C. All OSPF routers in totally stubby areas can have type 3 LSAs in their database.
D. All OSPF routers in totally stubby areas can have type 7 LSAs in their database.
E. All OSPF routers in NSSA areas can have type 3 LSAs in their database.
F. All OSPF routers in NSSA areas can have type 7 LSAs in their database.

Answer: AEF

QUESTION 50
When verifying OSPF virtual link problems, which is an important item to check on the two transit

OSPF routers? Select the best response.

A. OSPF process ID
B. OSPF router ID
C. OSPF network type
D. OSPF memory usage
E. OSPF CPU utilization
F. OSPF stub area configurations

Answer: B
Explanation:
The OSPF router IDs of the two transit OSPF routers are used to form the virtual link (with the area area-id virtual-link neighbor-router-id command) so it is an important item to check -> B is correct.

QUESTION 51
You are developing a verification plan for an upcoming OSPF implementation. Part of this plan is to verify the status of type 3 LSAs within the network. Which routers should you verify first to ensure that the configurations are correct for generating type 3 LSAs?
Select the best response.

A. Internal routers within the backbone area (area 0)
B. Internal routers within the NSSAs
C. Internal routers within the stubby areas
D. ASBRs
E. ABRs
F. DRs and BDRs

Answer: E
Explanation:
Type 3 LSA (Summary LSA) is advertised by the ABR of originating area to advertise network from other areas so we should check the ABRs first.

QUESTION 52
Which condition must be satisfied before an EIGRP neighbor can be considered a feasible successor? Select the best response.

A. The neighbor's advertised distance must be less than or equal to the feasible distance of the current successor.
B. The neighbor's advertised distance must be less than the feasible distance of the current successor.
C. The neighbor's advertised distance must be greater than the feasible distance of the current successor.
D. The neighbor's advertised distance must be equal to the feasible distance of the current successor.
E. The neighbor's advertised distance must be greater than or equal to the feasible distance of the current successor.

Answer: B
Explanation:
The feasible successor route is a route which has a higher metric than the successor route to reach a subnet but meets the feasibility condition and can be used in the event that the successor route goes down. This route does NOT get installed in the routing table but is kept in the topology table. The feasibility condition states that the AD from a neighbor must be less than the metric of

the successor route (the feasible distance [FD]) because routing through a feasible successor when the AD > FD may cause a routing loop.

QUESTION 53
Based on the need to limit processing and bandwidth utilization due to dynamic routing protocol operation, the following routing requirements have been specified for your network.
- partial and incremental routing updates
- only the devices affected by a topology change perform route recomputation
- route recomputation only occurs for routes that were affected
Which dynamic routing protocol should be deployed in your network to best meet these requirements? Select the best response.

A. BGP
B. OSPF
C. IS-IS
D. EIGRP
E. RIPv2

Answer: D
Explanation:
The bandwidth utilization issue has been addressed by implementing partial and incremental updates. Therefore, only when a topology change occurs does routing information get sent. Regarding processor utilization, the feasible successor technology greatly reduces the total processor utilization of an AS by requiring only the routers that were affected by a topology change to perform the route recomputation. Furthermore, the route recomputation only occurs for routes that were affected. Only those data structures are accessed and used. This greatly reduces search time in complex data structures.

QUESTION 54
Which statement about a non-zero value for the load metric (k2) for EIGRP is true? Select the best response.

A. A change in the load on an interface will cause EIGRP to recalculate the routing metrics and send a corresponding update out to each of its neighbors.
B. EIGRP calculates interface load as a 5-minute exponentially weighted average that is updated every 5 minutes.
C. EIGRP considers the load of an interface only when sending an update for some other reason.
D. A change in the load on an interface will cause EIGRP to recalculate and update the administrative distance for all routes learned on that interface.

Answer: C
Explanation:
The load metric (k2) represents the worst load on a link between source and destination. EIGRP routing updates are triggered only by a change in network topology (like links, interfaces go up/down, router added/removed), and not by change in interface load or reliability. The load is a five minute exponentially weighted average that is updated every five seconds (not five minutes) . EIGRP considers the load of an interface only when sending an update for some other reason (like a link failure)

QUESTION 55
Your network consists of a large hub-and-spoke Frame Relay network with a CIR of 56 kb/s for each spoke. Which statement about the selection of a dynamic protocol is true? Select the best

response.

A. EIGRP would be appropriate if LMI type ANSI is NOT used.
B. EIGRP would be appropriate, because the Frame Relay spokes could be segmented into their own areas.
C. EIGRP would be appropriate, because by default, queries are not propagated across the slow speed Frame Relay links.
D. EIGRP would be appropriate, because you can manage how much bandwidth is consumed over the Frame Relay interface.

Answer: D
Explanation:
By default, EIGRP will limit itself to using no more than 50% of the interface bandwidth. The primary benefit of controlling EIGRP's bandwidth usage is to avoid losing EIGRP packets, which could occur when EIGRP generates data faster than the interface line can absorb it. This is of particular benefit on Frame Relay networks, where the access interface bandwidth and the PVC capacity may be very different.

QUESTION 56
When an EIGRP topology change is detected, what is the correct order of events when there is a FS? Select the best response.

A. The neighbor adjacency is deleted.
The feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
B. DUAL is notified.
Remove all topology entries learned from that neighbor.
The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
C. The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
D. DUAL is notified.
The neighbor adjacency is deleted.
Remove all topology entries learned from
that neighbor. The feasible route is used.

Answer: D
Explanation:
If a packet is not received before the expiration of the hold time, the neighbor adjacency is deleted, and all topology table entries learned from that neighbor are removed, as if the neighbor had sent an update stating that all the routes are unreachable. If the neighbor is a successor for any destination networks, those networks are removed from the routing table, and alternative paths, if available, are computed. This lets the routes quickly reconverge if an alternative feasible route is available.

QUESTION 57
Refer to the exhibit. Why is the 140.140.0.0 network not used as the gateway of last resort even though it is configured first? Select the best response.
R3#show run | include default-

ip default-network 140.140.0.0
ip default-network 130.130.0.0
R3#show ip route | begin Gateway
Gateway of last resort is 0.0.0.0 to network 130.130.0.0
116.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
C 116.16.37.0/30 is directly connected, Serial1/0.2
C 116.16.32.0/30 is directly connected, Serial2/0.2
C 116.16.34.0/28 is directly connected, Serial1/0.1
C 116.16.35.0/28 is directly connected, Serial2/0.1
S 116.0.0.0/8 [1/0] via 116.16.34.0
* 140.140.0.0/32 is subnetted, 3 subnets
O 140.140.1.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
O 140.140.3.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
O 140.140.2.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1
* 130.130.0.0/16 is variably subnetted, 4 subnets, 2 masks
D* 130.130.0.0/16 is a summary, 00:30:04, Null0
C 130.130.1.0/24 is directly connected, Ethernet0/0
C 130.130.2.0/24 is directly connected, Ethernet0/1
C 130.130.3.0/24 is directly connected, Ethernet1/0
D 150.150.0.0/16 [90/679936] via 116.16.35.5, 00:02:58, Serial2/0.1

A. The last default-network statement will always be preferred.
B. A route to the 140.140.0.0 network does not exist in the routing table.
C. Default-network selection will always prefer the statement with the lowest IP address.
D. A router will load balance across multiple default-networks; repeatedly issuing the show ip route command would show the gateway of last resort changing between the two networks.

Answer: B
Explanation:
As you can see in the exhibit, 140.140.0.0 doesn't appear in the routing table.

QUESTION 58
Refer to the exhibit. Why are the EIGRP neighbors for this router not learning the routes redistributed from OSPF? Select the best response.

```
router eigrp 123 redistribute ospf 123
network 116.16.35.0 0.0.0.255 network 130.130.0.0
auto-summary
!
router ospf 123
log-adjacency-changes
network 116.16.34.0 0.0.0.255 area 0
neighbor 116.16.34.4
```

A. Redistribution must be enabled mutually (in both directions) to work correctly.
B. Auto-summary causes the OSPF routes redistributed into EIGRP to be summarized; thus the OSPF network 116.16.34 is summarized to 116.34.0.0, which is already covered by the EIGRP protocol.
C. Default metrics are not configured under EIGRP.
D. Both routing protocols must have unique autonomous system numbers for redistribution to function correctly.

Answer: C
Explanation:
Same as RIP, when redistribute into EIGRP from OSPF, the default metric is infinite -> We must set a seed metric when redistributing into EIGRP.

QUESTION 59
Which BGP option is required when load sharing over multiple equal-bandwidth parallel links from a single CE router to a single ISP router over eBGP? Select the best response.

A. eBGP Multipath
B. eBGP Multihop
C. BGP Synchronization
D. Public AS numbers

Answer: B

QUESTION 60
Which BGP feature should be used to avoid high memory utilization on a router? Select the best response.

A. soft-reconfiguration
B. route refresh
C. BGP communities
D. full-mesh BGP peering

Answer: B
Explanation:
BGP routers have enormous routing tables so it uses much memory to proceed these routes. When a BGP policy is changed, the BGP session needs to be reset for the policy to take effect. But the resetting results in route churn and route flapping. There are two ways to clear a BGP session without resetting the TCP session between them (this is often called "soft reset"):
Soft-reconfiguration: stores all received (inbound) routing policy updates without modification in a table so that when a new filter is applied, the router will use this table to calculate the changes without resetting the TCP session between the two BGP peers . This is a memory-intensive (high memory utilization) method and is not recommended.
Route-refresh: allows a BGP router to request a remote peer resend its BGP Adj-RIB-Out. This allows the BGP router to reapply the inbound policy. The route-refresh capability requires no extra memory on the local router